

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant:	Bruce Schofield	Examiner:	John Pauls
Serial No.:	10/805,993	Art Unit:	3686
Conf. No.:	1210		
Filed:	March 22, 2004	Attorney Docket No.:	909430-US-NP
Title:	METHOD AND APPARATUS FOR PROVIDING NETWORK BASED LOAD BALANCING OF MEDICAL IMAGE DATA		

CERTIFICATE OF ELECTRONIC TRANSMISSION

I hereby certify that this document, along with any other papers referred to as being attached or enclosed, is being filed electronically on June 11, 2010.

/John C. Gorecki/
John C. Gorecki, Reg. No. 38,471

M.S. Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF

(i) Real Party In Interest

This application is owned by Avaya, Inc., of Basking Ridge, New Jersey, USA.

(ii) Related Appeals and Interferences

None

(iii) Status of Claims

Claims 37-52 are pending in the application and stand rejected. Claims 1-36 have been canceled. The rejection of claims 37-52 under 35 USC 103 is appealed.

(iv) Status of Amendments

In response to the Final Office Action dated October 16, 2009, applicants filed a response on February 26, 2010, but did not propose any claim amendments. In an Advisory Action dated March 15, 2010, the Examiner indicated that the response had been considered. A pre-appeal brief request for review was then filed along with a notice of appeal on April 13, 2010. A

decision of the panel was entered May 3, 2010. No claim amendments were proposed after receipt of the Final Rejection. Accordingly, there are no un-entered amendments.

(v) Summary of Claimed Subject Matter

Medical imaging modalities may generate one image, a sequence of images, or a very large number of images. These images are frequently very high resolution, and may need to be processed in various ways such as to produce a three-dimensional volume image (Specification at paragraph 4). Frequently, all medical images are required to be handled by one or a small number of medical imaging archive systems such as a Picture Archive and Communication System (PACS), which can form a bottleneck on the network. (Specification at paragraph 5).

Applicants proposed to deploy a network service on the communication network that is configured to extract information from medical image data to determine what type of task will need to be performed by an image archive resource when the medical image data is provided to the image archive resource. The network service also monitors the available image archive resources. Using these two pieces of information – the complexity of the task required to be performed on the medical images and the available capacity of the image archive resources – the network service is able to load balance medical image processing tasks between the available image archive resources. (See e.g. Specification at paragraph 6).

The following tables map support in the specification to the claim limitations:

Independent Claim 37:

Claim limitation	Support
37. A method for performing network based load balancing of medical image data among a plurality of image archive resources by a network service deployed on a network element implemented on the network, the method comprising the steps of:	Fig. 1, Paragraphs 15-16: Network 104 Network element 106 Image Archive Resources 108a-108c Network service 112
monitoring, by the network service, a parameter associated with each of the plurality of image archive resources indicative of an available capacity of each of the plurality of image archive resources;	Paragraph 21; Fig. 2, reference 202

receiving, by the network service, medical image data having embedded therein instructions associated with a task to be performed by at least one of the image archive resources in connection with the medical image data;	Paragraph 22; Fig. 2, reference 204
extracting, by the network service, the instructions associated with the task from the medical image data;	Paragraph 22
determining, by the network service, a level of complexity of the task to be performed from the instructions associated with the task;	Paragraph 22; Fig. 2, reference 206
selecting, by the network service, one of the plurality of image archive resources to be used to perform the task in connection with the medical image data using, as a selection function, the available capacity of each of the plurality of image archive resources and the level of complexity of the task to be performed; and	Paragraph 23; Fig. 2, reference 208
transferring, by the network service, the medical image data to the selected one of the plurality of image archive resources.	Paragraph 24; Fig. 2, reference 210

Independent Claim 45:

Claim limitation	Support
45. An apparatus for performing network based load balancing of medical image data among a plurality of image archive resources, the apparatus comprising:	
a network element coupled to a network and configured to send and receive data via the network; and	Fig. 1, Paragraphs 15-16; Fig. 3, paragraphs 25-27: Network element 106

a network service coupled to the network element, the network service configured to:	Fig. 1, Paragraphs 15-16; Fig. 3, paragraphs 27-28: Network service 112
monitor a parameter associated with each of the plurality of image archive resources indicative of an available capacity of each of the plurality of image archive resources;	Paragraph 21; Fig. 2, reference 202
receive medical image data having embedded therein instructions associated with a task to be performed by at least one of the image archive resources in connection with the medical image data;	Paragraph 22; Fig. 2, reference 204
extract the instructions associated with the task from the medical image data;	Paragraph 22
determine a level of complexity of the task to be performed from the instructions associated with the task; and	Paragraph 22; Fig. 2, reference 206
select one of the plurality of image archive resources to be used to perform the task in connection with the medical image data using, as a selection function, the available capacity of each of the plurality of image archive resources and the level of complexity of the task to be performed; and	Paragraph 23; Fig. 2, reference 208
transfer the medical image data to the selected one of the plurality of image archive resources.	Paragraph 24; Fig. 2, reference 210

(vi) Grounds of Rejection to be Reviewed on Appeal.

Whether claims 37-42 and 45-50 are unpatentable under 35 U.S.C. 103 over Rothschild (U.S. Patent Application Publication No. 2002/0016718) in view of Primak (U.S. Patent No. 6,389,448), and Martin (U.S. Patent No. 6,263,368). Whether claims 43-44 and 51-52 are unpatentable under 35 U.S.C. 103 over Rothschild in view of Primak and Martin, and further in view of Carr (U.S. Patent No. 6,301,617) and Liu (U.S. Patent No. 5,031,089).

(vii) Argument

Claims 37-42 and 45-52

Rejection under 35 USC 103

Independent claims 37 and 45 were rejected under 35 USC 103 as unpatentable over Rothschild (U.S. Patent Application Publication No. 2002/0016718) in view of Primak (U.S. Patent No. 6,389,448), and Martin (U.S. Patent No. 6,263,368).

Independent claim 37 relates to a method for performing network based load balancing of medical image data among a plurality of image archive resources. Independent claim 45 relates to an apparatus for performing network based load balancing of medical image data among a plurality of image archive resources. Both claims recite that the level of complexity of the task to be performed should be used during the load balancing process.^{1 2} However, none of the cited references teach or suggest that the level of complexity of the task to be performed should be used in load balancing of medical images. Accordingly, the Examiner erred in rejecting independent claims 37 and 45 and the rejection under 35 USC 103 should be reversed.

¹ In claim 37, the method includes the steps of “determining, by the network service, a level of complexity of the task to be performed from the instructions associated with the task” and “selecting ... one of the plurality of image archive resources to be used to perform the task ... using, as a selection function, the available capacity of each of the plurality of image archive resources and the level of complexity of the task to be performed.”

² In claim 45, the apparatus includes a network service configured to “determine a level of complexity of the task to be performed...” and “select one of the plurality of image archive resources to be used to perform the task ... using, as a selection function, the available capacity of each of the plurality of image archive resources and the level of complexity of the task to be performed.”

In rejecting claim 37, (See Final Office Action dated October 16, 2009 at page 3) the Examiner cited Rothschild as teaching a network element and the method steps of (1) receiving medical image data by the network element and (2) transferring the medical image data to an image archive resource. The Examiner then cited Primak as teaching a network service and the method steps of: (3) monitoring the available capacity of each of the plurality of image archive resources; and (4) selecting one of the image archive resources using, as a selecting function, the available capacity of each of the plurality of image archive resources. (See Office Action at page 3-4).

On page 5 of the Office Action, (See Final Office Action dated October 16, 2009 at page 5), the Examiner acknowledges that the combination of Rothschild and Primak fails to teach or suggest the method steps of (5) determining, by the network service, a level of complexity of the task to be performed from the instructions associated with the task; and (6) selecting, by the network service, one of the plurality of image archive resources to be used to perform the task in connection with the medical image data using, as a selection function, the available capacity of each of the plurality of image archive resources and the level of complexity of the task to be performed. However, the Examiner contends that this is taught by Martin. Applicants respectfully traverse this position as not supported by the references.

For completeness, applicants have addressed each of the references cited by the Examiner. However, since the Examiner has admitted that the combination of Rothschild and Primak does not teach or suggest using the level of complexity of the task to be performed in connection with load balancing medical images to medical image archive resources, the real issue to be resolved is whether Martin teaches this feature. If Martin does not teach/suggest this feature, then the combination of Rothschild, Primak, and Martin collectively fails to teach/suggest this feature and, accordingly, the rejection of the independent claims should be reversed.

Rothschild

Rothschild teaches a medical image management system that may be implemented by an Application Service Provider to provide network based delivery and storage of medical images. The medical image management system allows users to have access to the medical images securely over the network and provides special clinical and visualization applications centrally

for the remote users. (Rothschild at paragraphs 136-140). Rothschild does not teach or suggest load-balancing images to image archives based on the complexity of the task required to be implemented on the medical images.

Primak

Fig. 2b of Primak shows a router 30 connecting three servers 10(a), 10(b) and 10(c) to the Internet. Each server has a Load Balancing module 12. A client passes a connection request (SYN packet) to the router which multicasts the connection request to all of the servers (Primak at Col. 3, lines 46-48: "After receiving the SYN packet from the client computer 60, the router 30 multicasts the SYN packet to all of the servers 10(a)-(c).").

In Primak, the load balancing module 12 of each server evaluates whether its associated server will handle the connection request by (1) calculating a pseudo-random number; and (2) determining the relative availability of each server. (Primak at Col. 3, lines 49-54: "The evaluation process involves calculating a pseudo-random number for each SYN packet and determining the relative availability of each server.") The load balancing module either passes the SYN packet (connection request) onto the server's stack or discards it based on its evaluation of the SYN packet. (Primak at Col. 3, lines 54-57). Each server generates the same pseudorandom number³ and has full information about each of the other servers' availability⁴ and, hence, each of the servers will make the same decision to implement distributed load balancing.

Primak does not teach or suggest that the complexity of the SYN packet or the task associated with the SYN packet should be used in making a load balancing decision when deciding which of the servers will process the SYN packet. Paragraph 22 of the instant specification specifies that "the level of complexity of a task will be based on the amount of resources that are required to execute the task, such as processor time or memory, or the

³ See Primak at Col. 3, line 57 to Col. 4, line 4: (explaining how the pseudo-random number is generated); see also Primak at Col. 4, lines 4-6: (explaining that each Load Balancing module generates an identical random number).

⁴ See Primak at Col. 4, line 8, "The relative availability of a server is a function of its overall capacity and current load."; see also Primak at Col. 4, lines 30-37: (teaching that an agent 14 on each server will periodically transmit the availability of the server to the other servers so that each of the load balancing agents knows how busy each of the other servers is).

complexity may be a function of the time that is required to execute the command.” These features look at the new task that is to be allocated, not the existing tasks that are already being processed by the various processors. Since the “task” in Primak is the SYN packet associated with establishment of a new connection, to find a corollary between paragraph 22 and Primak the Examiner would need to show that Primak looks at the complexity of the SYN packet. Primak does not do this.

In the Response to arguments section, (Final Office Action at page 12, lines 14-18), the Examiner asserted that Primak teaches assessing the complexity of a task to be performed in connection with load balancing, citing Primak at Col. 4, lines 30-37 and Fig. 2b. Specifically, the Examiner stated that Primak “discloses monitoring CPU capacity, CPU load, number of tasks being performed and the number of connections” which the Examiner stated “has the same meaning as complexity in the specification of the present application as shown in paragraph 22”. (Final Office Action dated October 16, 2009, at page 12, lines 16-18). As noted above, however, each of these factors looks at how the servers are operating, nothing in this list refers to the complexity of the new SYN packet that is to be assigned to one of the servers. Hence, the Examiner erred by equating monitoring server operation with evaluating the complexity of the task to be load balanced.

Martin

Martin teaches that conventional processor load balancing may break down where there is a large amount of data to be transmitted from a server. In this instance, the availability of bandwidth on the network, rather than processor bandwidth, may cause a bottleneck (Martin at col. 2, lines 61-67). Thus, Martin suggest that network server link loading should be monitored and used as the basis of performing load balancing between servers (Martin at Col. 3, 36-42).

The Examiner cited Martin at Col. 1, lines 41-44 Col. 3, lines 42-48, and Col. 5, lines 14-28 as teaching network load balancing based on the complexity of the task to be performed. In Col. 1, lines 41-44, Martin states that tasks should be distributed equally among the individual server computers to balance the overall loading of the server site to obtain optimum performance. At Col 3, lines 42-48, Martin states that a message traffic monitor is configured to monitor parameters representative of message traffic to and from the servers on the network server links. This reinforces the position outlined above, which is that Martin is looking at the

volume of traffic on the network to server links in connection with load balancing. This does not mean that Martin is looking at the complexity of the task to be performed but rather means that he is looking at the amount of traffic on the links connecting the servers to the network. At Col. 5, lines 14-28, Martin states that client requests are dispatched to servers by looking at “parameters representative of network loading on the server network links. (See Col. 5, lines 24-26).

Thus, Martin teaches that network to server “link loading” may be monitored and used as the basis of performing load balancing between servers (Martin at Col. 3, 36-42). Martin does not, however, teach or suggest that the load balancer should look at the complexity of the task to be allocated when selecting a server.

At the bottom of Page 5 of the Office Action (Final Office Action dated October 16, 2009, at page 5), the Examiner stated:

“Martin discloses that complexity is based on a number of criteria including: network loading, processor loading and packet or byte counts for example. These criteria are the same as, or are directly related to “complexity” as disclosed in the specification of the present application in paragraph 22.”

No independent basis for this assertion was provided in support of this assertion, however in the main portion of the rejection the Examiner cited to Martin at col. 3, lines 30-48 and col. 4, lines 3-5. In this section of Martin, Martin does discuss network loading, processor loading and packet or byte counts. However, as discussed below, Martin does so in evaluating network conditions or conditions of the processor, not in connection with evaluating a task to be load balanced. Accordingly, the Examiner’s conclusion that Martin discloses “complexity” is incorrect and not supported by the reference.

Specifically, the Examiner’s argument is not persuasive since it does not address the claim language. As noted above, Martin was cited by the Examiner as teaching two steps – “determining ... a level of complexity of the task to be performed from the instructions associated with the task”, and “selecting ... one of the plurality of image archive resources ... using, as a selection function, the available capacity of each of the plurality of image archive resources and the level of complexity of the task to be performed”. (emphasis added).

As is clear from this claim language, the level of complexity that is used in the load balancing process is first determined from instructions associated with the task. The criteria

cited by the Examiner “network loading, processor loading and packet or byte counts” are not quantities that are determined from instructions associated with a task to be load balanced. The “network loading,” for example, is a quantity that looks at how much traffic is present on the network. This is not extracted from instructions associated with the task to be load balanced, but rather looks at external network conditions.

Similarly, “processor loading” looks at how much capacity of a processor is in use at the time. This criterion is not extracted from instructions associated with the task to be load balanced, but rather looks at how the processor (which may ultimately handle the task) is running.

The third criterion mentioned by the Examiner “packet or byte counts” also is not extracted from instructions associated with the task to be load balanced, but rather is affiliated with how the load balancer determines the amount of traffic on the network links. Specifically, at Col. 3, lines 49-51, Martin states that “any load on the network can be measured.” Then, at Col. 3, lines 53-57, Martin explains that preferably the network load is measured by a message traffic monitor. In the paragraph spanning from Col. 3, line 62 to Col. 4, line 5, Martin explains how traffic monitor can use load counts to measure network load. As is clear from reading these portions of Martin in context, Martin looks at the load on the network by counting packets or bytes of traffic passing through the dispatcher or passing an external monitor. This is not information extracted from instructions associated with the task to be load balanced but rather is information that is unrelated to the task at hand and associated with other traffic on the network.

The Examiner stated that monitoring processor network loading, processor loading and packet or byte counts was “the same as, or directly related to ‘complexity’ as disclosed in the specification of the present application...” citing paragraph 22. For convenience, paragraph 22 has been reproduced below:

A task is received by the network service that is addressed to the image archive system for execution thereby (204). Typically, the task will be embedded within a DICOM message, which may also include image or textual data as well. The network service identifies the task within the DICOM message and extracts it. The network service determines the level of complexity of the received task, or assigns a priority to the task using predetermined criteria (206). Typically, the level of complexity of a task will be based on the amount of resources that are required to execute the task such as processor time or memory, or the complexity may be a function of the time that is required to execute the command.

Alternatively, the network service may assign a priority to the task. The priority may be based on predetermined criteria and provided to the network service or stored therein. The priority may be based on the source address of the message, the destination address of the message, the individual sending the message, or on the command itself.

As is clear from this paragraph, the network service identifies the task within the DICOM message and extracts it. The network service will then determine the complexity of the task which will be based: (1) on the amount of resources that are required to execute the task such as processor time or memory; or (2) as a function of the time that is required to execute the command. Both of these “complexity” items are based on an estimation of the extent and type of resources that will be needed to be expended to complete the task. The aspects identified by Martin are associated with resources that are already in use to process other transactions, and are not related to the complexity of the task that is to be load balanced. The Examiner’s statement that the teachings of Martin are the same as or directly related to the “complexity” recited in the claims is thus clearly incorrect.

Accordingly, applicants respectfully submit that the Examiner erred in rejecting the independent claims of this application and, accordingly, the rejection should be reversed. Specifically, no reference cited by the Examiner teaches or suggests a method of load balancing in which a network service performs the steps of (1) determining a level of complexity of the task to be performed from the instructions associated with the task; and (2) selecting one of the plurality of image archive resources to be used to perform the task in connection with the medical image data using, as a selection function, the available capacity of each of the plurality of image archive resources and the level of complexity of the task to be performed. The Examiner’s rejection, accordingly, should be reversed.

In the response to arguments section of the Office Action (See Final Office Action dated October 16, 2001, at Page 12, lines 8-12 and Page 12, lines 18-20), the Examiner twice stated that “one cannot show non-obviousness by attacking references individually where the rejections are based on combinations of references.” Applicants are not attacking references individually, but rather are asserting that none of the applied references teach or suggest a particular aspect of the claims. Specifically, none of the cited references teach or suggest that the level of complexity of the task to be performed should be used in load balancing of medical images.

Hence, if none of the references teach this claimed feature, the combination necessarily also does not teach or suggest the claimed feature. Accordingly, the Examiner committed legal error in rejecting the claims of this application and the rejection under 35 USC 103 should be reversed.

Claims 43-44 and 51-52

Claims 43-44 and 51-52 are dependent claims which were rejected over the same combination of references (Rothschild, Primak, and Martin) along with several secondary references. The secondary references do not make up the deficiencies noted above. Hence, for the reasons set forth above, applicants respectfully submit that the rejection of claims 43-44 and 51-52 should also be reversed.

(viii) Claims Appendix

An appendix containing the current version of all pending claims is attached.

(ix) Evidence Appendix

None.

(x) Related Proceedings Appendix

None

Conclusion

Applicants respectfully request that the rejections of claims 37-52 under 35 U.S.C. 103 be reversed.

If any fees are due in connection with this filing, the Commissioner is hereby authorized to charge payment of the fees associated with this communication or credit any overpayment to Deposit Account No. 502246 (Ref: 909430).

Respectfully Submitted

Dated: June 11, 2010

Anderson Gorecki & Manaras LLP
P.O. Box 553
Carlisle, MA 01741
Tel: (978) 264-4001
Fax: (978) 264-9119

/John C. Gorecki/
John C. Gorecki, Reg. No. 38,471

APPENDIX – PENDING CLAIMS

1-36. (Canceled)

37. A method for performing network based load balancing of medical image data among a plurality of image archive resources by a network service deployed on a network element implemented on the network, the method comprising the steps of:

monitoring, by the network service, a parameter associated with each of the plurality of image archive resources indicative of an available capacity of each of the plurality of image archive resources;

receiving, by the network service, medical image data having embedded therein instructions associated with a task to be performed by at least one of the image archive resources in connection with the medical image data;

extracting, by the network service, the instructions associated with the task from the medical image data;

determining, by the network service, a level of complexity of the task to be performed from the instructions associated with the task;

selecting, by the network service, one of the plurality of image archive resources to be used to perform the task in connection with the medical image data using, as a selection function, the available capacity of each of the plurality of image archive resources and the level of complexity of the task to be performed; and

transferring, by the network service, the medical image data to the selected one of the plurality of image archive resources.

38. The method of claim 37 wherein the step of selecting comprises selecting the one of the resources having the greatest available capacity relative to the complexity level of the task to be performed.

39. The method of claim 37, wherein the medical image data is formatted as a DICOM message.

40. The method of claim 37, wherein the plurality of image archive resources comprises a plurality of Picture Archive Systems (PACS).

41. The method of claim 40 wherein the parameter is one of the group consisting of the PACS server load, or the PACS storage time.

42. The method of claim 37 wherein, the parameter is one of the group consisting of the resource load, the capacity of the network, or the congestion of the network.

43. The method of claim 37, wherein the selection function is further based on a priority level of the task and a priority level of the resource.

44. The method of claim 43, wherein the step of selecting comprises selecting the one of the plurality of resources having a priority less than or equal to the priority level of the task.

45. An apparatus for performing network based load balancing of medical image data among a plurality of image archive resources, the apparatus comprising:

a network element coupled to a network and configured to send and receive data via the network; and

a network service coupled to the network element, the network service configured to:

monitor a parameter associated with each of the plurality of image archive resources indicative of an available capacity of each of the plurality of image archive resources;

receive medical image data having embedded therein instructions associated with a task to be performed by at least one of the image archive resources in connection with the medical image data;

extract the instructions associated with the task from the medical image data;

determine a level of complexity of the task to be performed from the instructions associated with the task; and

select one of the plurality of image archive resources to be used to perform the task in connection with the medical image data using, as a selection function, the

available capacity of each of the plurality of image archive resources and the level of complexity of the task to be performed; and

transfer the medical image data to the selected one of the plurality of image archive resources.

46. The apparatus of claim 45 wherein the network service is operative to select the one of the resources having the greatest available capacity relative to the complexity level of the task to be performed.

47. The apparatus of claim 45, wherein the medical image data is formatted as a DICOM message.

48. The apparatus of claim 45, wherein the plurality of image archive systems comprises a plurality of Picture Archive System (PACS).

49. The apparatus of claim 48 wherein the parameter is one of the group consisting of the PACS server load or the PACS storage time.

50. The apparatus of claim 45 wherein the parameter is one of the group consisting of the resource load, the capacity of the network, or the congestion of the network.

51. The apparatus of claim 45, wherein
the selection is further based on a priority level of the task and a priority level of the image archive resource.

52. The apparatus of claim 51, wherein the selection is based on selecting the one of the plurality of resources having a priority less than or equal to the priority level of the task.